

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (currently amended) A method for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the method comprising:

providing a shared bandwidth bucket associated with each of the plurality of the guaranteed bandwidth buckets;

allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in any one of the plurality of guaranteed bandwidth buckets;

determining whether bandwidth in one of the plurality of guaranteed bandwidth buckets is sufficient to allow traffic to pass immediately through the network appliance; and

transferring bandwidth from the shared bandwidth [[buck]] bucket to one of the plurality of guaranteed bandwidth buckets when it is determined that bandwidth in one of the plurality of guaranteed bandwidth buckets is not sufficient to allow traffic to pass immediately through the network appliance.

2. (original) The method of claim 1 wherein the shared bandwidth bucket is a token bucket.

3. (original) The method of claim 1 wherein the guaranteed bandwidth buckets are token buckets.

4. (original) The method of claim 1 wherein the guaranteed bandwidth buckets are credit/debit buckets.

5. (original) The method of claim 1 wherein each guaranteed bandwidth bucket is associated with a traffic shaping policy.

6. (original) The method of claim 1 wherein a plurality of guaranteed bandwidth buckets are associated with a single traffic shaping policy.

7. (original) The method of claim 5 wherein the traffic shaping policy screens based on IP address.

8. (previously presented) The method of claim 7 wherein the traffic shaping policy screens based on source IP address.

9. (previously presented) The method of claim 7 wherein the traffic shaping policy screens based on destination IP address.

10. (previously presented) The method of claim 7 wherein the traffic shaping policy screens based on protocol type.

11. (previously presented) The method of claim 7 wherein the traffic shaping policy screens based on UDP/TCP port number.

12. (original) The method of claim 7 wherein the traffic shaping policy screens based on the type of service requested.

13. (previously presented) The method of claim 5 wherein the traffic shaping policy screens based on traffic content.

14. (previously presented) A method for allocating bandwidth in a network appliance comprising:

defining a guaranteed bandwidth allocation for a first policy for passing traffic through the network appliance including using a first bucket to allocate the guaranteed bandwidth;

defining a guaranteed bandwidth allocation for a second policy for passing traffic through the network appliance including using a second bucket to allocate the guaranteed bandwidth;

sharing excess bandwidth developed from the underutilization of the guaranteed bandwidth allocated to the first and second buckets including

providing a shared bandwidth bucket associated with the first and second buckets; and

borrowing bandwidth from the shared bandwidth bucket by one of the first and second buckets when the respective bucket has insufficient bandwidth to allow traffic to pass immediately through the network appliance.

15. (original) An apparatus for allocating bandwidth in a network appliance where the network appliance includes a plurality of guaranteed bandwidth buckets used to evaluate when to pass traffic through the network appliance, the apparatus comprising:

a shared bandwidth bucket associated with a plurality of the guaranteed bandwidth buckets;

means for allocating bandwidth to the shared bandwidth bucket based on the underutilization of bandwidth in the plurality of guaranteed bandwidth buckets; and

a scheduler operable to
evaluate a packet to determine if a traffic shaping policy should be applied to a given packet,

evaluate a guaranteed bandwidth bucket associated with an identified traffic shaping policy,

determine when the guaranteed bandwidth bucket associated with an identified traffic shaping policy has insufficient capacity to support a transfer of the packet through the network, and

borrow bandwidth from the shared bandwidth bucket by a respective guaranteed bandwidth bucket to allow traffic to pass immediately through the network appliance.

16. (previously presented) A network device comprising:

a first bucket configured to receive tokens at a first information rate;

a second bucket configured to receive tokens at a second information rate;

a third bucket configured to receive extra tokens from the second bucket; and

a scheduler configured to:

determine if a size of traffic received at the network device exceeds a number of tokens stored in the first bucket,

determine, when the size of the traffic does not exceed the number of tokens stored in the first bucket, if a size of the traffic exceeds a number of tokens stored in the second bucket, and

transfer, when the size of the traffic exceeds the number of tokens stored in the second bucket, an appropriate number of tokens from the third bucket to the second

bucket so that the second bucket includes a number of tokens that equals or exceeds the size of the traffic.

17. (previously presented) The network device of claim 16 wherein the scheduler is further configured to:

cause the traffic to be forwarded after the transfer; and

decrement the number of tokens in the first and second buckets based on the size of the traffic.

18. (previously presented) The network device of claim 16 wherein the scheduler is further configured to:

determine if the third bucket includes the appropriate number of tokens, and

prohibit the traffic from being forwarded when the third bucket includes less than the appropriate number of tokens.

19. (previously presented) The network device of claim 16 further comprising:
one or more input ports configured to receive traffic from a network, each of the one or more input ports including the first bucket, the second bucket, the third bucket, and the scheduler.

20. (previously presented) A method comprising:

receiving traffic;

determining if a policy is to be applied to the traffic;

determining, when a policy is to be applied to the traffic, if a size of the traffic exceeds a number of tokens in a first bucket, the first bucket being associated with the policy;

determining, when the size of the traffic does not exceed the number of tokens in the first bucket, if the size of the traffic exceeds the number of tokens in a second bucket;

determining, when the size of the traffic exceeds the number of tokens in the second bucket, if a third bucket includes an appropriate number of tokens that, when added to the number of tokens in the second bucket, would equal or exceed the size of the traffic;

transferring the appropriate number of tokens from the third bucket to the second bucket when the third bucket includes the appropriate number of tokens; and

forwarding the traffic after the transferring.

21. (previously presented) The method of claim 20 further comprising:

forwarding the traffic when the size of the traffic does not exceed the number of tokens in the second bucket.

22. (previously presented) The method of claim 20 further comprising:

repeating the determining if a size of the traffic exceeds a number of tokens in a first bucket; the determining, when the size of the traffic does not exceed the

number of tokens in the first bucket, if the size of the traffic exceeds the number of tokens in the second bucket;

the determining, when the size of the traffic exceeds the number of tokens in the second bucket, if a third bucket includes an appropriate number of tokens that, when added to the number of tokens in the second bucket, would equal or exceed the size of the traffic; and the transferring the appropriate number of tokens from the second bucket to the first bucket when the third bucket includes the appropriate number of tokens for at least a second policy prior to transferring the traffic, the second policy being associated with different first and second buckets.